In the Specification

Please replace the title appearing at page 1, lines 1-2 with the following:

SYSTEM AND METHOD FOR DISTRIBUTING PACKETS AMONG A PLURALITY OF PATHS TO A DESTINATION

Please replace the paragraph beginning at page 7, line 22 with the following:

FIG. 1 shows a simplified, high-level block diagram that schematically illustrates network 10 as associated with exemplary embodiments of the present invention. Network routing devices 20A, 20B, 20C, 20D, 20E are coupled together via the exemplary paths 32, 34, 36,38, 40, 42. For example, path 32 is coupled from port 26 of network device 20A to a port of network device 20B, and path 34 coupled from port 30 of network device 20A to a port of network device 20D. Similarly, path 36 is coupled between network devices 20B and 20E, path 38 between network devices 20C and 20E, and path 40 between network device 20D and 20E. Accordingly, a plurality of paths 32-36, or 42-38 or ~~34,40~~ 34-40 are available for propagating data from a source 22 to destination 24.

Please replace the paragraph beginning at page 8, line 17 with the following:

Referencing FIG. 2, in simplified, high-level block diagram overview, a network routing device 20 comprises packet ~~forwarding circuitry~~ processors 44xx to interface various parts of the router as associated with exemplary embodiments of the present invention. Line cards 52, 54, 56 and 58 provide physical ports 78, 80, 82 and 84 of the routing device 20, with line card 52, for example, providing up to 24 gigabit Ethernet ports 78 and line card 54 providing similar capabilities. Line card 56 includes two 10 gigabit Ethernet ports 82, while line card 58 provides an OC-192 POS (packet-over-SONET) port 84. Although the four line cards are shown, many systems provide slots to accommodate additional cards, for example, of up to 14 line cards in a particular exemplary embodiment. Accordingly, a user can configure routing device 20 to accommodate different traffic models and physical port mixes by appropriate selection of line cards and programming.

Please replace the paragraph beginning at page 8, line 29 with the following:

Switching-fabric 50 moves each data packet from an ingress port of a line card to an egress port of a line card. Switching-fabric 50 connects to each line card through two switching-fabric port connections (see, for example, port connections 62, 64 to line card 52, port connections 66, 68 to line card 54, port connections 70, 72 to line card 56, and port connections 74, 76 to line card 58). Switching-fabric

50 can be rapidly reconfigured. For example, at one instant, it may be switching packets from ingress port 68 to egress port 70 and at another instant from ingress port 64 to egress port 76, and at the next instant, it could be switching packets from ingress port 62 to egress port 66. At the same time, packets from ingress port 70 could be routed to egress port 68, and packets from ingress port 74 could be routed to egress port 62.

Please replace the paragraph beginning at page 9, line 26 with the following:

Referencing FIGS. 3-4, an exemplary packet processor 44 for designating or identifying egress ports includes processor 110A to receive data 113 at its input 114. Packet processor 44 uses a simple destination based next-hop method 442, shown in flowchart 440 of Figure 4, to establish an egress port ID for a packet. Processor 110A determines ~~144~~ 444 a destination address DA from the header of the received data and uses the destination address 112 as a basis to index ~~146~~ 446 content addressable memory (CAM) 120. Upon matching an index 121A of its table 122A, content addressable memory 120 will output next-hop pointer 123A at its pointer output 124. The next-hop pointer output indexes ~~148~~ 448 next-hop table 130, which comprises a lookup table 132, having a plurality of port identifications 133 selectably indexed by respective next-hop pointers 131. Upon establishing a pointer match to a particular index 131 therein, the next-hop table 130 outputs an egress port identification 133 at its output 134.

Please replace the paragraph beginning at page 10, line 27 with the following:

Referencing FIGS. 5-6, in accordance with an exemplary embodiment of the present invention, a plurality of equal cost paths interface respective plurality of egress ports 26, 28 and 30 of the network-processing device 20A for distributing data flows from ingress port 22 to destination 24. A processor 110B of packet processor 44B obtains HASH information from a header of a received data packet 113. Packet processor 44B uses a destination-plus-hash based next-hop method 652, shown in flowchart 650 of Figure 6, to establish an egress port ID for a packet. For example, the additional HASH information can be obtained from various portions of the header, such as, destination address, source address, data length, IP protocol, control information, preamble, or other equivalent regions per an IP protocol. According to a particular exemplary embodiment, a destination address is determined and then combined with the additional hash information ~~154~~ 654 to provide an index 112 for indexing ~~156~~ 656 a content addressable memory (CAM) 120B. CAM 120B is programmed as lookup table 122B having a plurality of entries 125. The entries include next-hop pointer values 123 indexed by respective index values 121. Distinguishing the CAM table 122A of FIG. 3, table 122B of the present exemplary embodiment includes multiple next-hop pointers 123 that are associated with a common destination address, for example, DA1;

yet, the separate entries are further distinguishable per the additional hash information. With processor 110B deriving index values 112 from both the destination address and the additional hash information as obtained from the header of the received data 113 and with CAM table 122B further configured with the additional next-hop pointers per destination address, the packet processor is able to distribute the data flows across a plurality of paths via the egress ports established by configuration table 122B, even though the flows may be destined for the same destination address.

Please replace the paragraph beginning at page 11, line 18 with the following:

As similarly noted before, the pointer 124 output from CAM 120B is used to index ~~148~~ 648 next-hop table 130 to match an index value 131 of its lookup table 132. Upon finding an index match, next-hop table 130 presents an egress port identification 133 at its output 134. Using this output 134 of the packet processor 44B, a router 20 (referencing FIG. 1 and 2) is able to transfer the received data to the identified egress port. Again, data-handling circuitry of the router can package, encapsulate or switch labels of the received data with appropriate new labels per its IP environment.

Please replace the paragraph beginning at page 13, line 17 with the following:

In accordance with a particular exemplary embodiment of the present invention, LAG-size table 150 determines a LAG-size representative of the number of links associated with an identified link aggregation LAG. If the LAG-size table comprises 32 entries of the different LAG IDs and each entry has a possibility of sixteen different LAG-size values, then there exists a total of 512 possible physical links that could be separately identifiable. As a part of determining which of the 512 possible links to use, for example, the LAG size value from the output 152 of the LAG-size table ~~100~~150 is forwarded to a LAG-distribution circuit, which comprises incrementor 154 and modulo processor 160. The distribution circuit selects a random integer of value between zero and sixteen, i.e., a value up to the number of links of the identified link aggregation. The random value that is provided from output 162 of the distribution circuit is combined with the LAG-ID 144 at combination node 164 to generate a next-hop pointer of value between zero and 512. The random value provided by modulo processor 160 establishes the least significant bits of the next-hop pointer 164 while the LAG ID provides the most significant bits. For example, assume the distribution circuit comprises four output bits 162, these output bits are then combined with five bits 144 of the LAG-ID at summary node 163 to provide a total of nine bits 164 of a next-hop pointer for indexing the next-hop table 170. The generated next-hop pointer 164 is used to index the next-hop table 170.

Please replace the paragraph beginning at page 14, line 16 with the following:

In accordance with an exemplary method 202 of distributing data flows across multiple links of a link aggregation, with reference to flowchart 200 of FIG.8, a data packet is received 204 and a destination address 206 determined from a header of the received data packet. Optionally, the method continues along path 230 to query whether or not a LAG operation is to be used, wherein the LAG operation query determination is based upon the destination address and the IP protocol of the data propagation. If no LAG operation is to be used, the method follows path 236, for determining 238 an identification of the next-hop egress port directly via the destination address. On the other hand, if query 232 determines that a LAG operation is to be used, then the method traverses alternative path 234, i.e., to paths 210,212.

Please replace the paragraph beginning at page 16, line 6 with the following:

When an entry of the second plurality of entries 155 of lookup table ~~142~~142B is indexed, the second field ~~133~~ 139 provides a LAG identification value less than 512, wherein the most significant bits of the second field are zero while the least significant bits, e.g., bits[0:8], of the second field provide a LAG identification. A portion (e.g., least significant bits) of the second output field 145 (FIG. 9A) of second table 140B is directed to establish a LAG-ID pointer 151 to index the LAG-size table 150. The first portion is additionally forwarded by line ~~144~~ 164 to grouping node163. LAG-ID pointer 151, taken from the first subportion, indexes LAG-size table 150 to obtain a number of links associated with the given identified LAG. The determined LAG-size is output to the distribution circuit 154, 160. The distribution circuit, similarly as presented before relative to FIG. 7, utilizes the LAG-size in combination with the hash value to determine, using a random distribution procedure, which one of the available number of links of the identified LAG to use in a next-hop data transfer.

Please replace the paragraph beginning at page 16, line 32 with the following:

Continuing with reference to FIGS. 9A and 9B, the second subportion (for example, the most significant bits) of the second output field 145 of second table 140B, serve to establish a control signal 252 for operating multiplexer 251. In accordance with a particular exemplary embodiment, when the second output field 145 of the second table 140 provides an output value less than 512, then logic 250 establishes a control signal 252 to configure multiplexer 251 to select next-hop pointer 175 for indexing next-hop table ~~176B~~ 172, 174 from the output 164 of the LAG circuit. Alternatively, when the second output field 145 from second table 140B provides a value greater than or equal to 512, then logic 250 establishes a control signal 252 to configure multiplexer 251 to select the next-hop pointer for indexing table 176B directly from the second output field 145.

Please replace the paragraph beginning at page 17, line 16 with the following:

Further referencing FIGS. 9A and 9B, when CAM 120 indexes entries 157 of the look-up table ~~142~~142B of the second table 140B, for example, with index values of PNTR5,PNTR6, . . . then next-hop pointer values of the second field 139 of magnitude greater than 512 are presented to the second output field 145 of second table 140B. Additionally, the third field 141 will present an active enable level signal (EN) to the third output field 148. The active level enable signal drives multiplexer 180 to select a next-hop egress-port ID from the output 176 of the next-hop table ~~176B~~ 176.

Please replace the paragraph beginning at page 17, line 24 with the following:

Next-hop table ~~176B~~ 172, 174 has a first plurality 172 of entries that comprise next-hop port ID for identifying ports to links of various link aggregations. A second portion 174 of the next-hop table ~~176B~~ 172, 174 includes port identifications to be indexed directly, via pointers of value greater than 512, by the second output field 145 of an external SRAM, or second table 140B.

Please replace the paragraph beginning at page 18, line 19 with the following:

Further referencing FIGS. 10A and 10B, second table 140C comprises a memory map or look-up table 142C having a plurality of fields: field-1, field-2, field-3, and index column 139. In a first plurality of entries 153' within the look-up table 142C (for example, of index values PNTR1, PNTR2, PNTR3 . . . of the index column 139), field-1 143 comprises a plurality of port identifications (e.g., egress-port-1, egress-port-2, egress-port-3) available for output to the first output field 146 of second table 140C when indexed appropriately by an index of CAM 120. Additionally, field-3 141 of lookup table 142C includes a (equal-cost-multi-path) count value equal to zero (CNT=0) to indicate the absence of an equal-cost, multi-path distribution. NOR-logic circuit 270 receives the ECMP count of the third output field 149 of second table 140C to establish enable signal 272 for controlling multiplexer 180. When the control signal input 272 is high, e.g., multiplexer 180 selects and egress-port identification as provided by the first output field of the second table 140C. Alternatively, when the control input 272 is low, multiplexer 180 selects ~~and~~ an egress-port identification from the output 176 of next-hop table 170C.

Please replace the paragraph beginning at page 19, line 7 with the following:

Continuing with reference to FIGS. 10A and 10B, when CAM 120 indexes an entry of the look-up table ~~142~~ 142C having a count value in the third field 141 equal to one, for example, entries indexed by PNTR5-PNTR6 . . ., then the first subportion SUB1 of the second field provides a next-hop pointer

(e.g., NHP1, NHP2 . . .) for presentation at output A of the second output field 143 of second table 140C. Additionally, modulo processor 264 will receive the count value of one from the third output field ~~145~~ 141 of table 140C and establish a control signal, e.g., of a zero level, at control input 262 to configure multiplexer 260 for selecting its first input A. Accordingly, the next-hop pointer of the first subportion A of the second output field 143 is used to index next-hop table 170C.

Please replace the paragraph beginning at page 20, line 27 with the following:

When CAM table 120B is selected, a pointer from CAM 120B indexes the alternative SRAM 140C-2, which will provide a plurality of data output fields 146,143,149 in place of the original SRAM 140C. Under this condition, the original SRAM may have its outputs tri-stated, while the alternative SRAM 140C-2 uses its respective first, second and third output fields to drive (not shown): (i) the first input of multiplexer 180, (ii) ports A . . . H. of multiplexer 260, and (iii) the count-value input of distribution circuit ~~164~~ 264 and the inputs of NOR-gate 270. In accordance with such alternative exemplary embodiment, more than eight equal-cost, multi-path pointer selections are selectively available for indexing next-hop table 170C.

Please replace the paragraph beginning at page 23, line 28 with the following:

It will be noted that the present exemplary embodiment of FIG. 13 can be viewed as a "layering" of two ECMP selections that allow for an increased number of selectable paths. Compare this to the exemplary embodiment described earlier herein with reference FIG. 10A, which uses a second CAM ~~130B~~ 120B to provide for additional plurality of selectable equal cost paths.